Bayesian games

Games and equilibrium concepts under uncertainty

Tomas Singliar

CS3150, April 21

Examples

- Blackjack
  - Action space = draw card/pass
  - Add up card values, maximum total < 22 wins
  - **What's the next card?**

- JPMorgan buys Bear Stearns
  - Action space = share price offered
  - Payoff = true value of business – payment
  - **What is the extent of the mess?**

- eBay
  - **How much did that other guy bid?**

Outline

Goal: to show how to incorporate incomplete information into games, define equilibrium concepts and show how tricky they become.

- Setup sequential games of incomplete information
- Sender-Receiver Games
- Decision making under uncertainty
- Bayesian Nash concepts and desirable properties

Sequential games

- Games where players take turns
- Optimal algorithm: decision tree
  - Zero-sum: “minimax tree”

Deterministic game

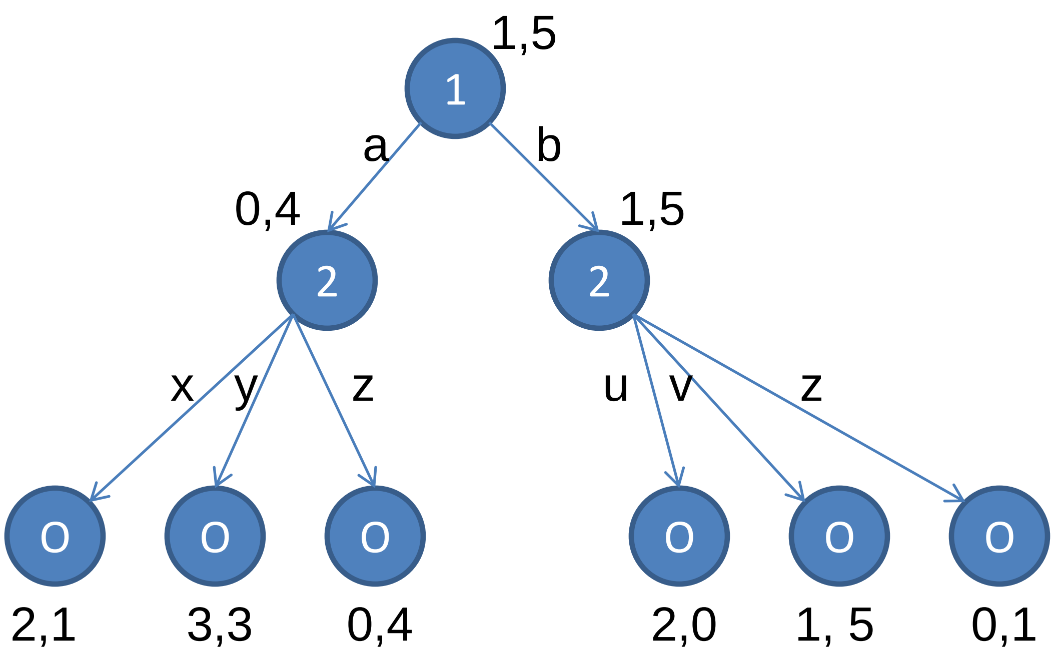

1 chooses max utility action

2 chooses max utility action

2’s actions conditioned on 1

Outcomes

Utilities of outcome (u_1,u_2)

Bayesian game setup

- Dynamic game – played sequentially
- Each player *i* has a set of ***types*** Θ_i

1. Nature selects player's **type** θ_i € Θ_i according to **p** – prior joint distribution over types, publicly known

- Player is **privately** informed of his type

Uncertainty = *players don't know other players' types*

2. Players now choose their actions a in some order

3. Players receive payoffs according to outcome

Payoff function on the outcome

$$u_i(a, \theta_i) \rightarrow R$$

Sender-Receiver Games

- Simplest case: two players – Sender, Receiver
  - **Running example**: Sender applies for job at Receiver. Her message is what she did last year, i.e. went to Beach or College. Receiver decides whether to Hire or Reject Sender. Type is fixed, i.e. going to College does not make you smart.
- Sender
  - Types: (private) θ € Θ = *{Smart, Dull}* **type space**
  - Action: send a message *m* € *M = {Beach, College}* **message space**
- Receiver
  - only one type (deterministic)
  - Action: *a* € *A = {Hire or Reject}*
  - Holds a prior belief $p(\theta)$ *over the Sender's type:* **action space**

    $p(\theta = Smart) = \gamma$ $p(\theta = Dull) = 1\text{-}\gamma$

This belief is assumed to be common knowledge.

Graphical representation

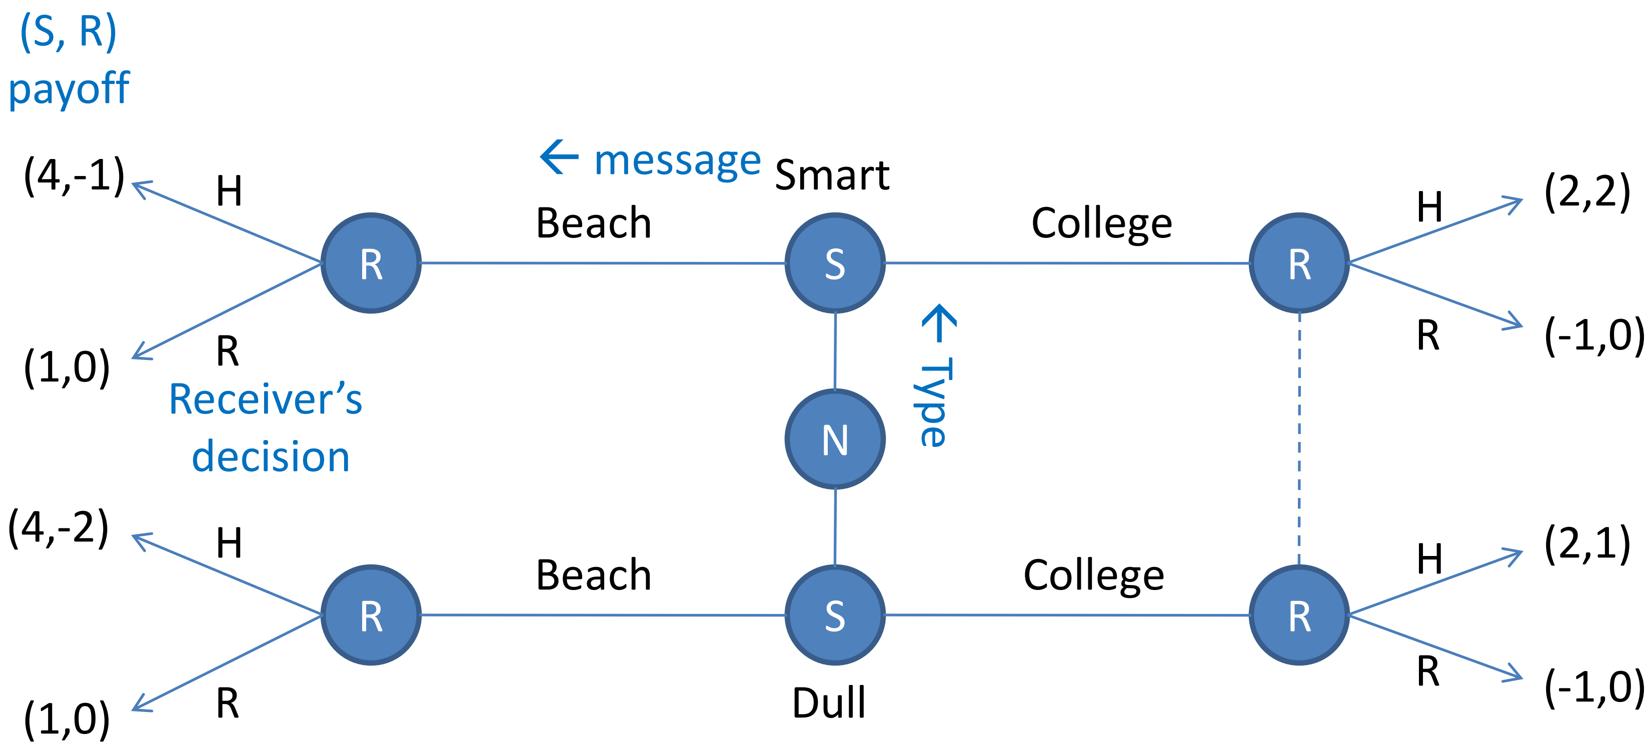

3 players: Nature, Sender, Receiver
2 sender types: Smart, Dull
1 receiver type

R does not know which part of the tree he is in. His **information set** = all nodes he could be in.

Graphical representation

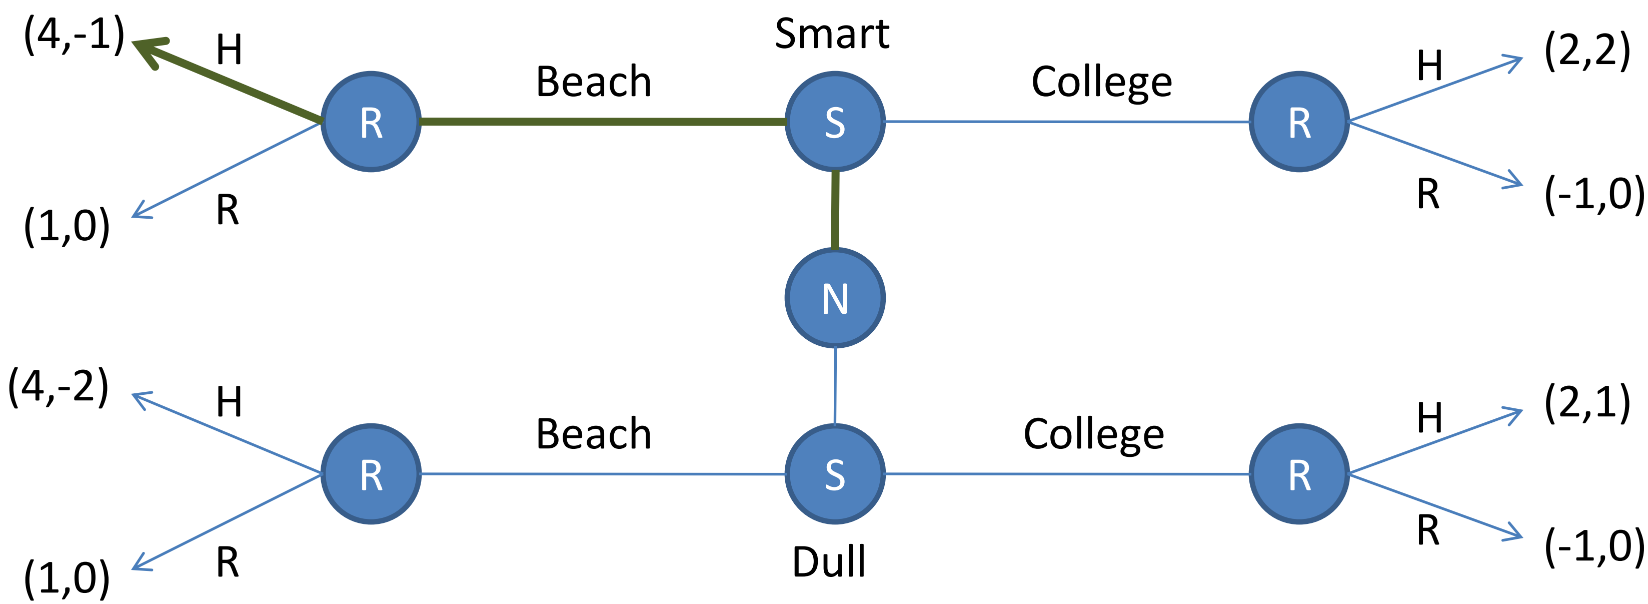

- Everything else being equal:
- Sender prefers to be Hired (by 3), go to the Beach (by 2). It doesn't matter if she's smart or dull.
- Receiver wants an educated applicant**, so Hires iff Sender went to College**

More notation

- Payoffs: *u* for Sender, *v* for Receiver

$$u, v : M \times A \times \Theta \to \mathscr{R}$$

- Pure strategies
  - Sender has only her type θ: $\bar{m} : \Theta \to M$
  - Receiver has only the message: $\bar{a} : M \to A$
- Mixed strategies
  - Sender – Prob. distribution over M $\sigma : \Theta \to \mathscr{M}$
  - Receiver – distribution over A $\rho : M \to \mathscr{A}$

Sender's best response

- Fix a message *m*, we know ρ, θ
- Payoff for m is expected value over R's reaction

$$E_a[u] = \sum_{a \in A} \rho(a \mid m) u(m, a, \theta)$$

- Sender's best responses then are the max-utility messages

$$\tilde{M}(\rho, \theta) =_{def} \arg\max_{m \in M} \sum_{a \in A} \rho(a \mid m) u(m, a, \theta)$$

- σ is a best response to ρ

  if it is nonzero only in $\tilde{M}(\rho, \theta)$, i.e.

$$\text{supp}\, \sigma(\theta) \subseteq \tilde{M}(\rho, \theta)$$

Receiver's best response

- Receiver has obtained *m* $\rho : M \to \mathscr{A}$
  - Knows $\sigma(\theta)$, but does not know θ
  - Let $\rho(a \mid m)$ be the prob. that Rcvr plays *a* after *m*
- Wishes to maximize **his** expected utility
- R updates its belief about S's type given *m*
  – *Message carries* ***a signal*** *about type of S*
- R decides based on the posterior belief

Receiver belief update

- Bayes rule

$$\tilde{p}(\theta \mid m) = \frac{p(\theta)\sigma(m \mid \theta)}{p(m)} = \frac{p(\theta)\sigma(m \mid \theta)}{\sum_{\theta' \in \Theta} p(\theta')\sigma(m \mid \theta')}$$

- If denominator is non-0, the message is **on-path** – some sender type has a non-zero probability of sending *m*

$$p(A, B) = p(A)\,p(B \mid A) \qquad p(A \mid B) = \frac{p(A, B)}{p(B)} = \frac{p(B \mid A)\,p(A)}{p(B)}$$

Receiver's decision

- Receiver maximizes his utility for each message separately

$$\tilde{A}(\tilde{p}, m) = \arg\max_{a \in A} \sum_{\theta \in \Theta} \tilde{p}(\theta \mid m) v(m, a, \theta)$$

- Strategy ρ is a best response to σ iff

$$\text{supp } \rho(m) \subseteq \tilde{A}(\tilde{p}, m)$$

- Remember, $\tilde{p}$ depends on σ

Bayesian equilibrium

Definition. A Bayesian equilibrium of the S-R game is a triple $(\sigma, \rho, \tilde{p}) \in \mathscr{M}^{\Theta} \times \mathscr{A}^{M} \times [\Delta(\Theta)]^{M}$ such that

- For all $\theta \in \Theta$, $\operatorname{supp} \sigma(\theta) \subseteq \tilde{M}(\rho, \theta)$
- For all on-the-path messages $m \in M^{+}(\sigma)$,
  $$\operatorname{supp} \rho(m) \subseteq \tilde{A}(\tilde{p}, m)$$
- The conditional posterior belief system $\tilde{p}$ is consistent with the Bayes rule

In our example…

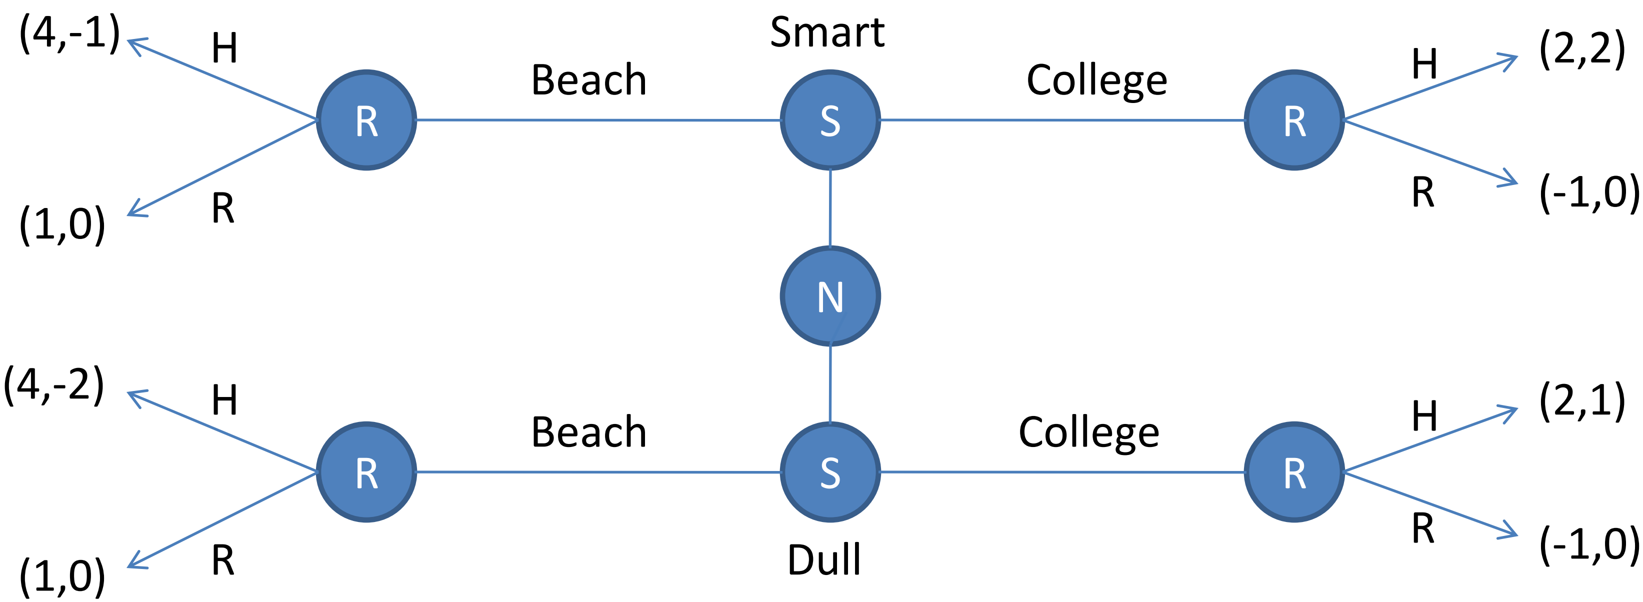

$(X, \quad Y, \quad L, \quad R, \quad s, \quad t)$

Sender action if Smart

Sender action if Dull

Receiver action
if Beach
if College

Receiver belief
of Smart
if Beach
if College

In our example… Eq 1

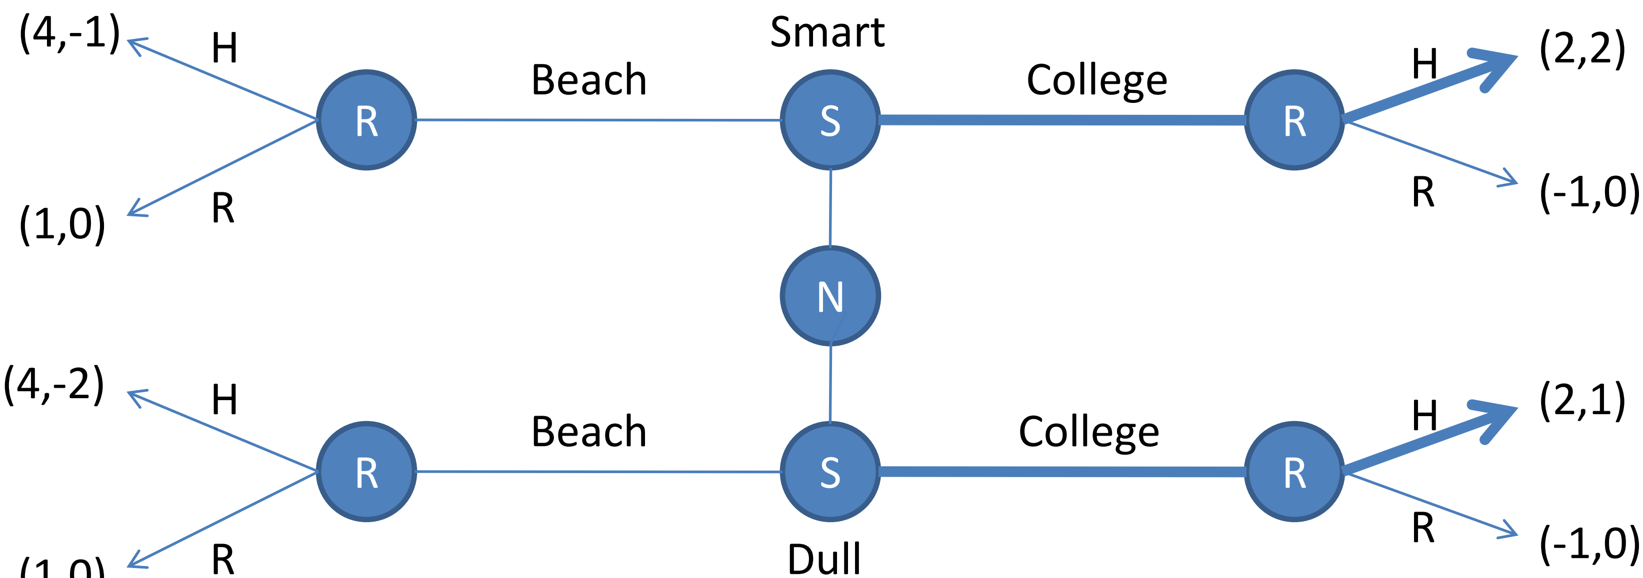

(C, C, R, H, *, γ)

Sender action if Smart

Sender action if Dull

Receiver action
if Beach
if College

Receiver belief
of Smart
if Beach
if College

This is a Bayesian equilibrium

Pooling strategy profile

In our example… Eq 2

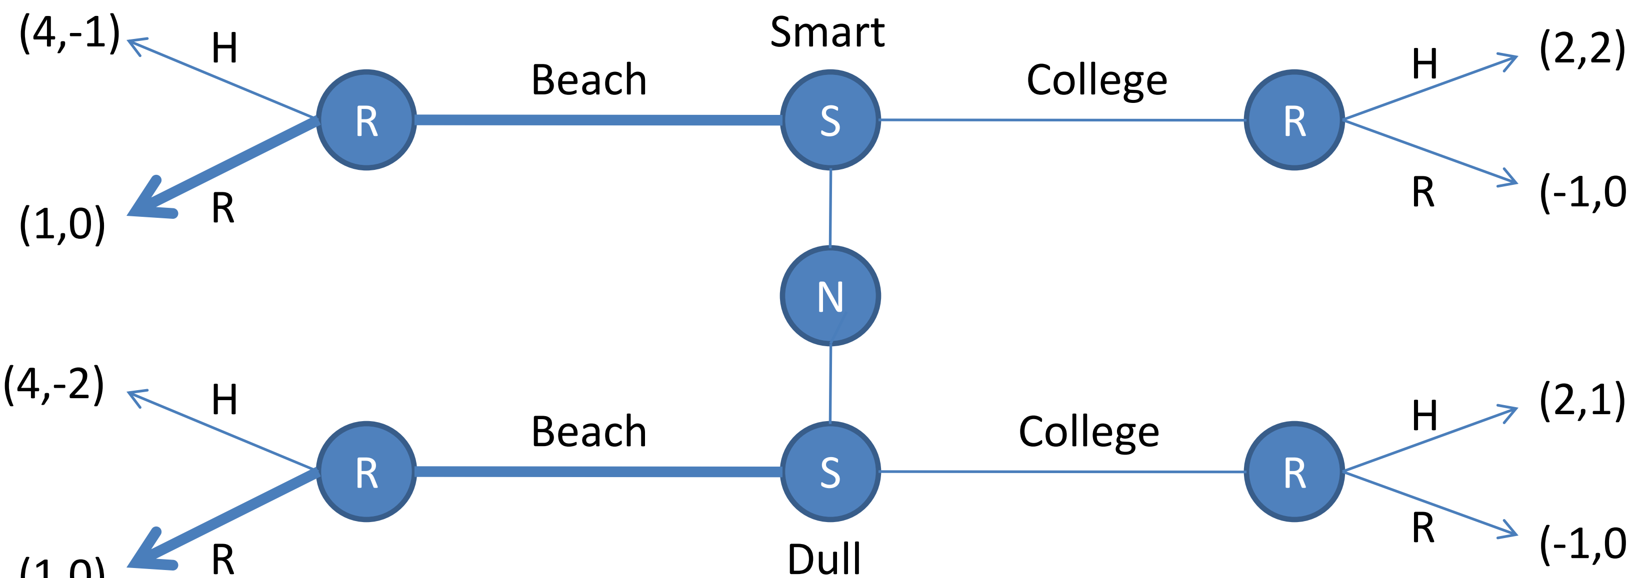

$$(B, \quad B, \quad R, \quad R, \quad \gamma, \quad *)$$

Sender action if Smart

Sender action if Dull

Receiver action
if Beach
if College

Receiver belief
of Smart
if Beach
if College

This is a Bayesian equilibrium **But an unstable one**

Perfect Bayesian equilibrium

Definition. A Bayesian equilibrium of the S-R game is a triple $(\sigma, \rho, \tilde{p}) \in \mathscr{M}^{\Theta} \times \mathscr{A}^{M} \times [\Delta(\Theta)]^{M}$ such that

- For all $\theta \in \Theta$, $\operatorname{supp} \sigma(\theta) \subseteq \tilde{M}(\rho, \theta)$
- For **all** ~~on-the-path~~ messages $m \in M^{+}(\sigma)$,

  $$\operatorname{supp} \rho(m) \subseteq \tilde{A}(\tilde{p}, m)$$
- The conditional posterior belief system $\tilde{p}$ is consistent with the Bayes rule

Existence of Bayesian equilibria

- **Nash thm**: A mixed Nash always exists if there are finitely many players with finite action sets.
- Create a regular game with $|\Theta|N$ players (replicate player for each type) $\{p_1^1, p_1^2, \ldots, p_1^{|\Theta_1|}, p_2^1, \ldots p_2^{|\Theta_2|}, \ldots\}$
  - With utilities: $v_i^j(s) = EU_i(s(\theta_i), s_{-i}(.); \theta_i^j)$

"Agent plays what the corresponding typed agent would in the corresponding normal game"

Nash thm gives an equilibrium mixture over strategies σ. Then $\varphi_i(\theta_j) = s'$ w.p. $\propto \sigma_i^j(s')$ is a Bayesian equilibrium.

More trouble with Bayesian equilibria

Additional conditions of equilibrium reasonableness

College is tough (for dulls)

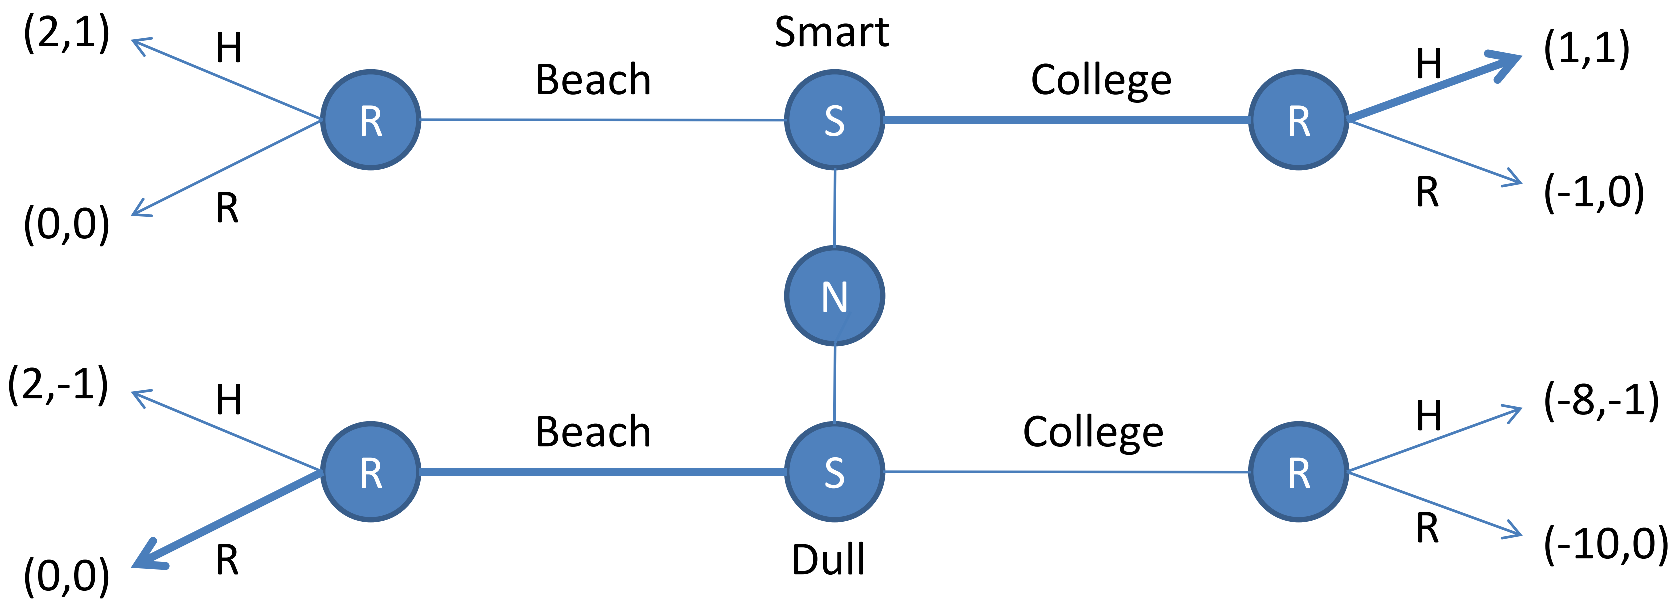

- Education is unproductive (no gain for Rcvr from Coll)
- But **signals** the type of applicant
- (C, B; R, H; 0, 1) is a **separating** PBE

Smart to College
Reject Beachgoers
Beachgoers are Smart wp 0

College is tough (for dulls)

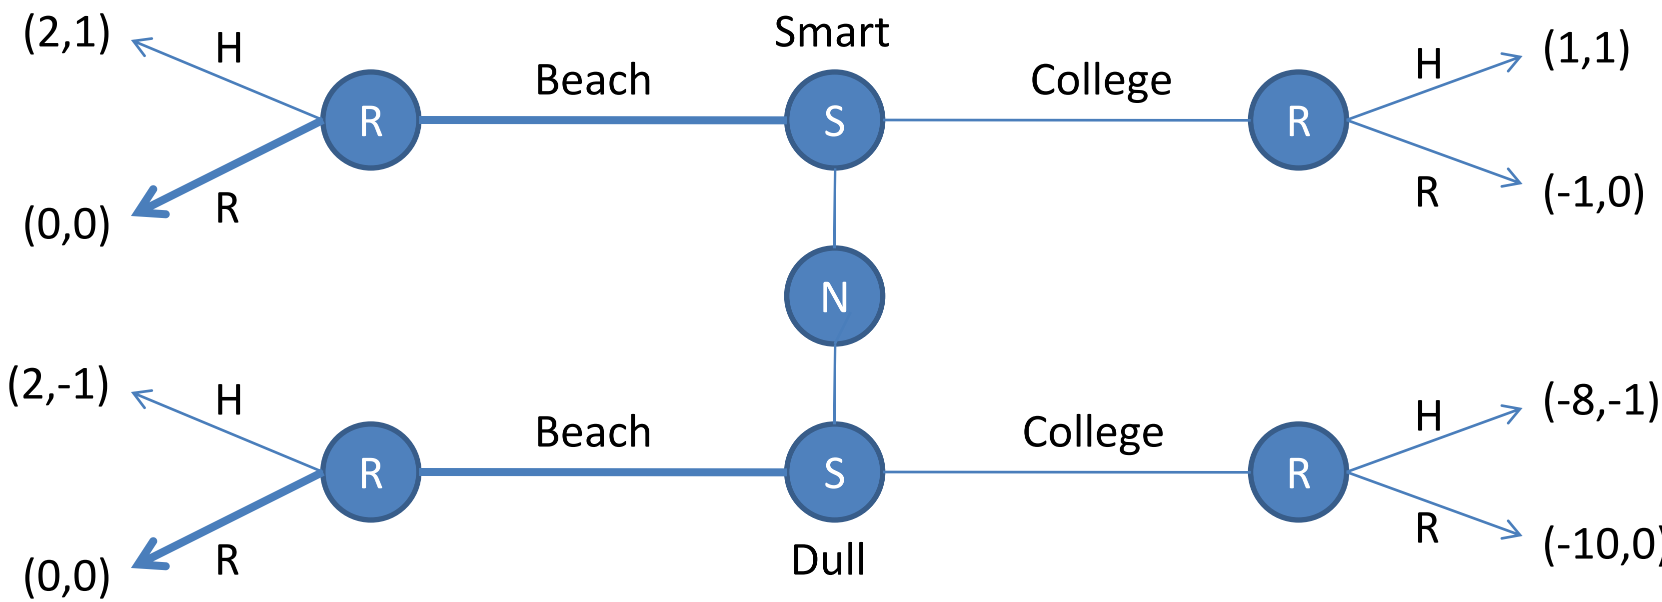

- Education is unproductive (no gain for Rcvr from Coll)
- (B, B; R, R; s, t), s < ½, t< ½ is a PBE
- Problem: Rcvr interprets deviation (C) as coming from a type (θ=D) who has no incentive to deviate

Everybody to Beach
Reject everybody
Beachgoers are Smart wp s

Dominated messages

- College is ***dominated*** for dulls: whatever the outcome (R or H), Dull type is better off with Beach

Definition. A message m is dominated for $\theta \in \Theta$ if there exists *m'* such that

$$\min_{a \in \hat{A}(m')} u(m', a, \theta) > \max_{a \in \hat{A}(m)} u(m, a, \theta)$$

where $\hat{A}(m)$ are all the actions that can be a best response (for some type $\theta \in \Theta$):

$$\hat{A}(m) = \bigcup_{p \in [\Delta(\Theta)]^M} \widetilde{A}(\widetilde{p}, m)$$

Test of dominated messages

A PBE fails the test of dominated messages if for any arbitrary message *m,* the Receiver puts a positive weight on the Sender being of the type for which the message is dominated.

== the Receiver assumes the Sender is irrational,

because she should have sent the dominat**ing** message

Technical note: there must be some type for which *m* is not dominated, otherwise a) technical problem b) we are computing a response to something a fully rational agent could not do.

Equilibrium domination

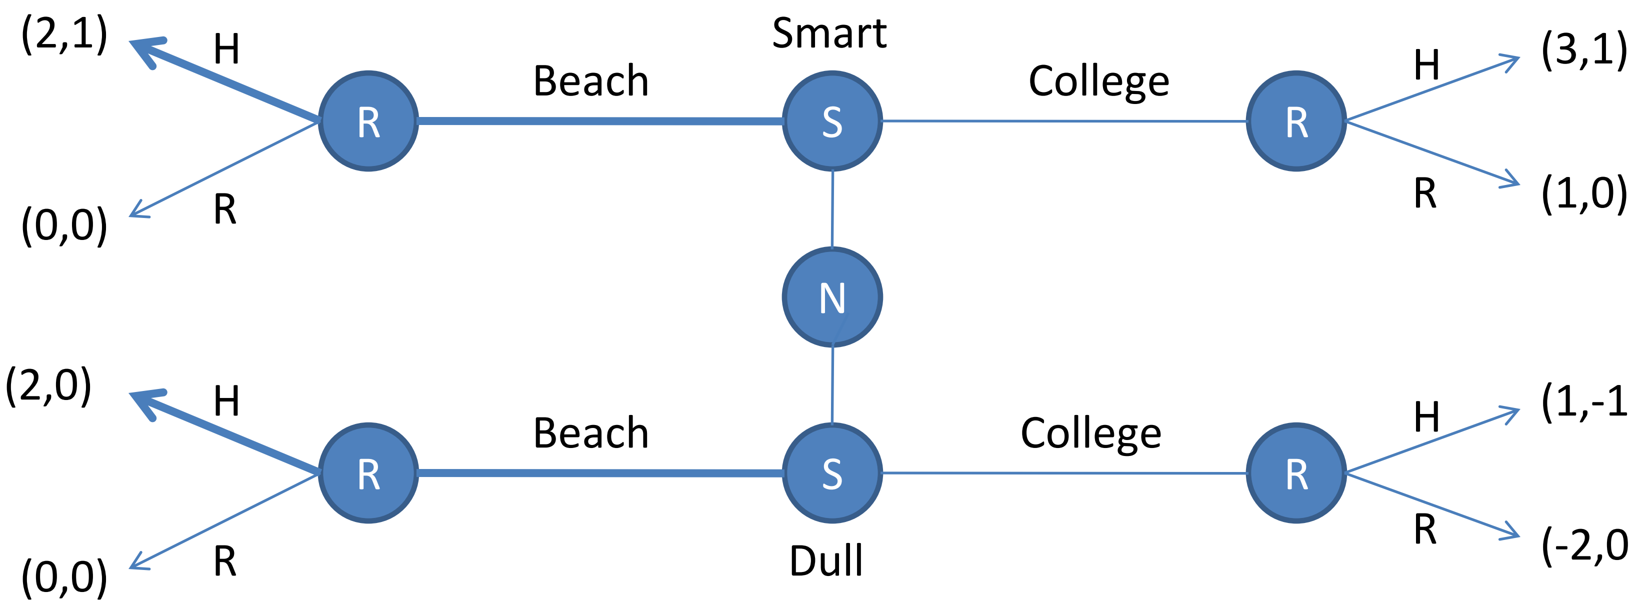

- (B, B; H, R; s, t), t< ½ is a BE that passes test of dominated messages
- Everybody gets hired for 2 pts but Smarts want to deviate

Equilibrium domination

- M is equilibrium dominated wrt eq ψ if expected payoff $\overline{u}(\theta)$ from the equilibrium exceeds what the player can get by playing m:

$$\overline{u}(\theta) > \max_{a \in \hat{A}(m)} u(m, a, \theta)$$

- The Dull sender should not deviate:
  – p(Dull|College) = 0

Thanks

Further readings

- Book Section 9.6 (Bayesian mechanism design)

Junk slides